arXiv:1401.2965v1 [cs.DC] 13 Jan 2014

A Hypermedia Distributed Application for Monitoring and Fault-Injection in Embedded Fault-Tolerant Parallel Programs

V. De Florio G. Deconinck M. Truyens W. Rosseel R. Lauwereins

Katholieke Universiteit Leuven
Electrical Engineering Dept. – ACCA
Kard. Mercierlaan 94 – B-3001 Heverlee – Belgium

Abstract

We describe a distributed, multimedia application which is being developed in the framework of the ESPRIT-IV Project 21012 EFTOS (Embedded Fault-Tolerant Supercomputing). The application dynamically sets up a hierarchy of HTML pages reflecting the current status of an EFTOS-compliant dependable application running on a Parsytec CC system. These pages are fed to a World-Wide Web browser playing the role of a hypermedia monitor. The adopted approach allows the user to concentrate on the high-level aspects of his/her application so to quickly assess the quality of its current fault-tolerance design. This view of the system lends itself well for being coupled with a tool to interactively inject software faults in the user application; this tool is currently under development.

1 Introduction

As systems get more and more complex, the need for a one-look snapshot of their activity is indeed ever increasing. This need has been strongly felt by people involved in the development of the ESPRIT-IV Project EFTOS [10] (Embedded Fault-Tolerant Supercomputing), whose aim is to set up a software framework for integrating fault-tolerance into embedded distributed high-performance applications in a flexible and easy way.

Through the adoption of the EFTOS framework, the target application running on a parallel computer is plugged into a hierarchical, layered system whose structure and basic components are:

- at the lowest level, a set of parametrisable functions managing error detection (Dtools) and error recovery (Rtools). A typical Dtool is a watchdog timer thread or a trap-handling mechanism; a Rtool is e.g., a fast-reboot thread capable of restarting a single node or a set of nodes. These are the basic components that are plugged into the embedded application to make it more dependable. EFTOS supplies a number of these Dtools and Rtools, plus an API for incorporating user-defined EFTOS-compliant tools;
- at the middle level, a distributed application called *DIR net* (detection, isolation, and recovery network) is available to coherently combine Dtools and Rtools, to ensure consistent strategies throughout the whole system, and to play the role of a backbone handling information to and from the fault-tolerance elements. To fulfill these requirements, the DIR net makes use of processes called Manager, Agents, and Backup Agents;
- at the highest level, these elements are combined into dependable mechanisms e.g., methods to guarantee fault-tolerant communication, voting methods and so on.

During the lifetime of the application, this framework guards the application from a series of possible deviations from the expected activity; this is done by executing detection, isolation, and reconfiguration tasks. For instance, a protection violation caught in a thread by a trap handling Dtool may trigger a relocation of that thread elsewhere in the system. As another example, if an error is detected which affects a component of the DIR net itself, say a Manager, then the system will isolate that component and elect another one (actually, one of the Backup Agents) as the DIR Manager.

To let the user keep track of events like those sketched above, the DIR net continuously prints on the system console short textual descriptions. Evidently such a linear, unstructured listing of events

pertaining different aspects of different actions taking place in different points of the user application, do not make up the best mechanism to gain insight in the overall state of the fault-tolerant system. On the contrary, a hierarchical, dynamic view of the structure and behavior of this system, including:

- its current shape (on which node which components are running, and their topology),
- the current state of its components (for instance, whether they are regarded to be correct, faulty, or are being recovered),
- each component's running history,

appeared to be the best solution fulfilling our needs.

Two main advantages from the adoption of such a strategy were foreseen, namely:

- (at design and system validation time) the possibility to assist the user assessing and/or validating his/her EFTOS-based fault-tolerance design,
- (at run time) the possibility to shorten the latency between the occurrence of the event, its comprehension, and a proper reaction at user level[1].

This work describes the architectural solution that has been successfully adopted within EFTOS to easily and quickly develop a tool fulfilling the above stated needs—a portable, highly customizable hypermedia monitor for the EFTOS applications making use of cheap, ready available off-the-shelf software components like e.g., the Netscape Navigator. It also shows how this monitor supplies the user with the needed structured information, and how it proved its usefulness within EFTOS. In particular an extension is described, currently under development, by means of which our monitor is turned into a versatile tool for fault-injection.

2 Design Requirements

In order to quickly deliver human-comprehensible information from the gigantic data stream produced by an EFTOS application, two needs have been assessed:

- a hierarchical representation of the data.

  (Most of the produced data is available, but it has to be organized and made browsable in "layers":
  - at the highest level, only the logical structure of the application should be displayed: which nodes are used, the EFTOS components executed on each node, and their overall state;
  - at a medium level, a concise description of the events pertaining each particular node should be made available;
  - at the lowest level, a deeper description of each particular event may also be supplied on user-demand),
- the use of multimedia.

  "An image is worth a thousand words", they say, and maybe even more insight can be derived from the extensive use of colors, sounds, video-clips and so on. For instance, re-coloring a green image to red may lead the user into realizing that a previously good situation has turned into a problem. The use of colors traditionally associated to meanings, or whose meaning can be borrowed from well-known everyday situations (e.g., those of traffic-lights) further speeds up the delivery of the information to the user.

Both things are available nowadays in products like Netscape or similar World-Wide Web [4] browsers which are able to render hierarchies of dynamically produced HTML [5] pages. We therefore decided to develop a distributed application piloting a WWW browser which in turn plays the role of a hypermedia renderer for the EFTOS system activity. This product, which we call the EFTOS Monitor, is described in the following Sections.

3 The Architecture of the EFTOS Monitor

The EFTOS Monitor basically consists of three components (see Fig. 1):

1. a *client* module, to be run by the DIR net;
2. an "*intermediate*" module, to be run by a number of Common Gateway Interface [15] (CGI) scripts;
3. the "*renderer*" i.e., a World-Wide Web browser.

- The client part, together with the DIR net and the user application, runs on a Parsytec CC system [1], a distributed-memory MIMD supercomputer consisting of powerful processing nodes based on PowerPC 604 at 133MHz, dedicated high-speed links, I/O modules and routers.

[1] Indeed, the high volume of data coming out of such a complex system is very likely to at least delay the appearance of the failure in the so-called *user's universe* i.e., "where the user of a system ultimately sees the effect of faults and errors" [14]; in some cases it may also make it transparent to the user altogether.

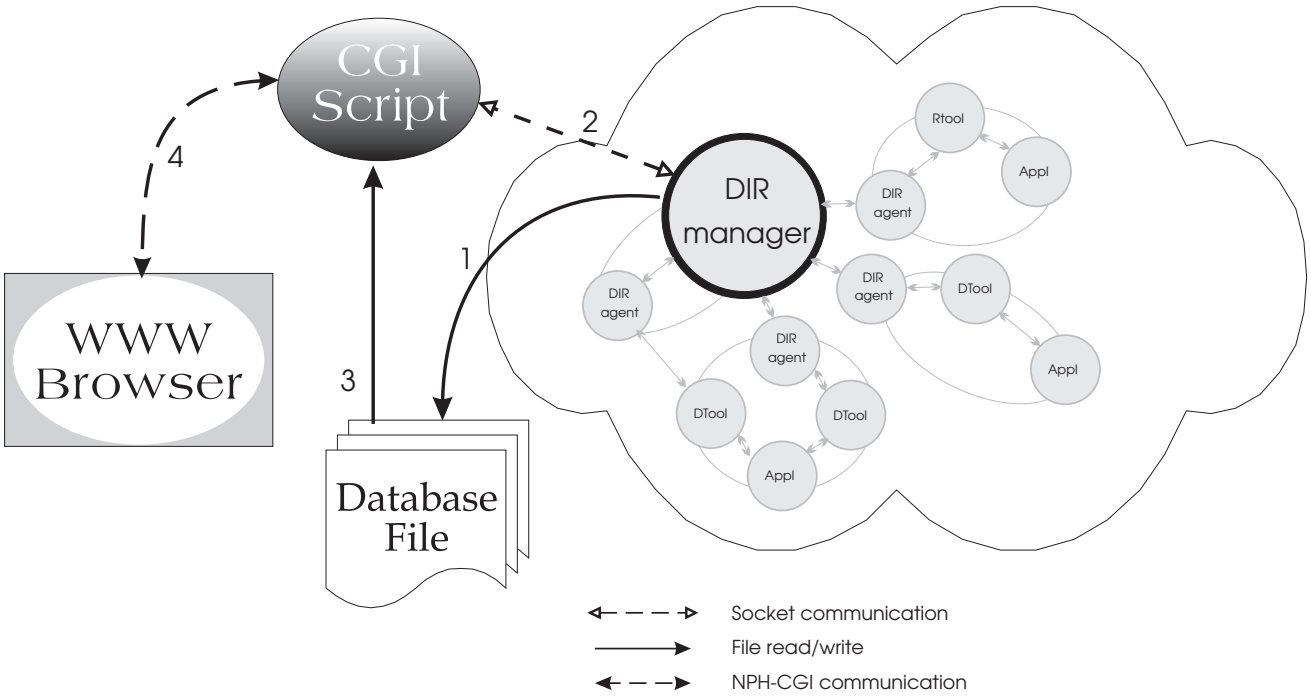

Figure 1: The architecture of the EFTOS Monitor: the CGI scripts and the EFTOS application share the same file system and communicate via a socket stream: each time a new event takes place, the DIR net updates a special database (1) and sends a notification to the main CGI script (2). The latter reads the database (3) on the arrival of the notification and converts it into a HTML hypertext, which is then fed (4) to a Netscape or another World-Wide Web browser for hypermedia rendering. The client part of the Monitor is integrated in the DIR net Manager process.

The system adopts the thread processing model; threads exchange messages through a proprietary message passing library called EPX [2] (*Embedded Parallel extensions to uniX*). The main tasks of the client module are the set up and the management of a database maintaining an up-to-date snapshot of the system activity, including the current mapping of the DIR net's components onto the processing nodes and the state and current activity of each component. This module also connects to the intermediate part via TCP sockets (see for instance [9]) and signals it on the very beginning and on the occurrence of each state transition.

- The intermediate module consists of a hierarchy of CGI scripts spawned by an Apache HTTP [6, 13] daemon, all running on the workstation hosting the CC system. The root script of this hierarchy connects with the client module and acts as a TCP server: for each new stimulus, the snapshot file is read over and a HTML document is produced and fed to the renderer. A connection is also started up with this latter so to be able to tightly interact with it without the intervention of the HTTP daemon: having done like this, one CGI script may stay alive and produce multiple HTML requests, which is not the case in ordinary CGI script—this special feature is known as "non-parsing header" (NPH) mode [15]. Logically speaking, we may say that the intermediate module acts as a gateway between the CC system and the hypermedia renderer. Like mythical Janus (It. *Giano*), one face is turned to the client module and gathers its requests, while the other is turned to the renderer and translates those requests in HTML—its main component has therefore been called `cgiano`.
- The third component, the renderer, simply is a browser like Netscape playing the role of a server able to display HTML data.

The application is started in two steps via a shell script whose first task is to run the renderer (or to

reconnect to a previously run renderer: this latter is possible using e.g., the remote control extensions [16] of Netscape, or an approach based on the Common Client Interface [12] mechanism of Mosaic; see for instance [11]). The renderer is run with a uniform resource locator [3] (URL) pointing to the root-level CGI script, which connects to Netscape in NPH-mode and starts listening for a TCP connection.

As a second step, the shell script spawns the parallel application on the CC system. Then the application launches the DIR net and the Monitor's client module; this latter initializes the snapshot files, connects to the CGI script and sends it the first signal. The script reacts to that stimulus by translating the main snapshot file in HTML and requesting the renderer to display it. The top-left image in Fig. 2 shows a typical output of this phase: the EFTOS application appears to the user as an HTML table depicting the processing nodes in the user partition. The state of each module is illustrated by means of colors with obvious meaning (green is "OK", red is "not OK", yellow means that the module is currently being recovered, and so on). In this way the user can immediately perceive whether a node is ready or not and which actions are carried out on it, as asked for in the requirements (§2.)

Information displayed in this HTML document only covers the logical structure of the application. If the user "clicks" any icon on this page, a high-level hypertextual description of the DIR net-events pertaining that specific node is displayed in a separate, cloned Netscape session (see Fig. 2, Window "Node-specific Information".) To keep this page up to date, an automatic reload is periodically performed. This technique is explained e.g., in [15].

This secondary document is in turn a hypertext whose links point to in-depth descriptions of each specific event (see Fig. 2, Window "Attached Information.")

4 Architecture Assessment

A number of observations may be drawn upon the above presented architecture; in particular:

- in our experience the architecture is easy and fast to design and develop, and effective especially towards fast prototyping;
- it is based on unmodified, low-cost, off-the-shelf hypermedia components which are widely available, continuously supported and updated on a wide range of hardware architectures;
- it is open, in the sense that the architecture is based on wide-spread standards e.g., the use of uniform resource locators [3] within a World-Wide Web interconnection, the HTML language, TCP/IP sockets, the MIME classification, and so on;
- it is distributed, and in particular the renderer may run on any X11-compliant Display Server, including a remote PC.

A possible alternative is to develop a custom application to play as a tailored monitoring tool. As an example, Scientific Computing Associates' TupleScope visual debugger is a custom X-based visualisation and debugging tool for parallel programs using the LINDA approach [8]. This may result in higher performance and possibly be more flexible but of course:

- it reasonably requires more time to develop even for a simple prototype;
- it requires custom design and development choices that may impact portability and supported features e.g., which software development environment and specifically which language and which libraries to use, or whether to restrict the hypermedia rendering to images or to use sounds as well—these choices may be simply skipped in our approach;
- distribution and hypermedia issues call for specific support which turn into higher costs and longer times.

For instance, TupleScope runs with the user application by adding a special linking option at compile time to the user application; this means it has been developed on purpose as a custom X11 application. Though it perfectly addresses its own goals, it has limited rendering capabilities (it only deals with static images) and would certainly require non-negligible efforts to adapt it towards other media. Moreover, TupleScope is available on a number of platform, though the costs of this portability and consequent support are certainly not negligible as well.

5 A Tool for Fault Injection

The same approach used to monitor the state of an EFTOS-compliant application is also effective in order to actively interact with it. Considering once again

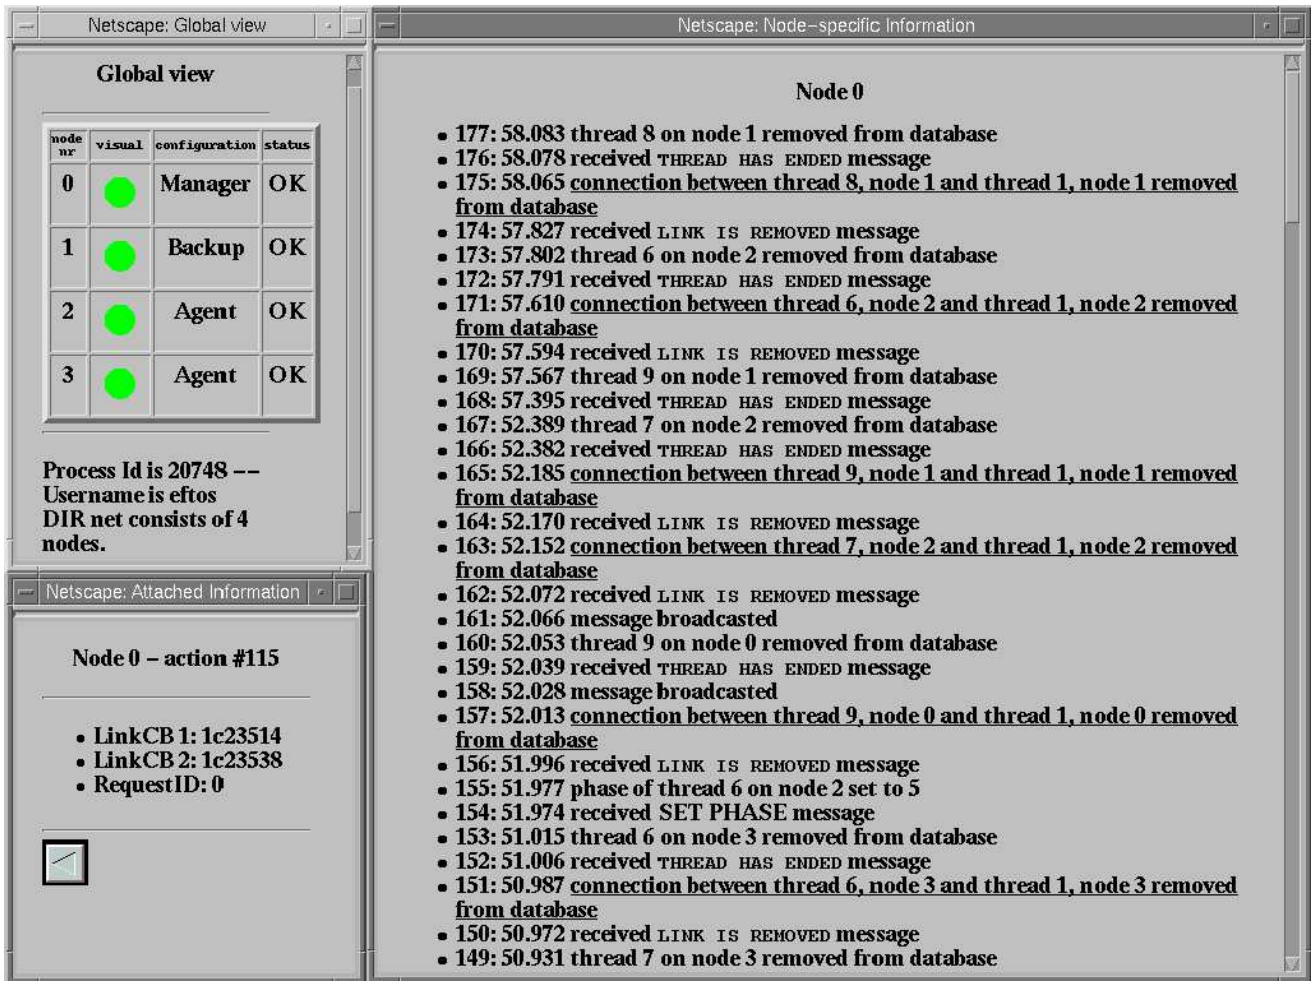
Netscape: Global view
Global view
node nr
visual
configuration
status
0
Manager
OK
1
Backup
OK
2
Agent
OK
3
Agent
OK
Process Id is 20748 --
Username is eftos
DIR net consists of 4 nodes.
Netscape: Attached Information
Node 0 – action #115
LinkCB 1: 1c23514
LinkCB 2: 1c23538
RequestID: 0
Netscape: Node-specific Information
Node 0
177: 58.083 thread 8 on node 1 removed from database
176: 58.078 received THREAD HAS ENDED message
175: 58.065 connection between thread 8, node 1 and thread 1, node 1 removed from database
174: 57.827 received LINK IS REMOVED message
173: 57.802 thread 6 on node 2 removed from database
172: 57.791 received THREAD HAS ENDED message
171: 57.610 connection between thread 6, node 2 and thread 1, node 2 removed from database
170: 57.594 received LINK IS REMOVED message
169: 57.567 thread 9 on node 1 removed from database
168: 57.395 received THREAD HAS ENDED message
167: 52.389 thread 7 on node 2 removed from database
166: 52.382 received THREAD HAS ENDED message
165: 52.185 connection between thread 9, node 1 and thread 1, node 1 removed from database
164: 52.170 received LINK IS REMOVED message
163: 52.152 connection between thread 7, node 2 and thread 1, node 2 removed from database
162: 52.072 received LINK IS REMOVED message
161: 52.066 message broadcasted
160: 52.053 thread 9 on node 0 removed from database
159: 52.039 received THREAD HAS ENDED message
158: 52.028 message broadcasted
157: 52.013 connection between thread 9, node 0 and thread 1, node 0 removed from database
156: 51.996 received LINK IS REMOVED message
155: 51.977 phase of thread 6 on node 2 set to 5
154: 51.974 received SET PHASE message
153: 51.015 thread 6 on node 3 removed from database
152: 51.006 received THREAD HAS ENDED message
151: 50.987 connection between thread 6, node 3 and thread 1, node 3 removed from database
150: 50.972 received LINK IS REMOVED message
149: 50.931 thread 7 on node 3 removed from database

Figure 2: The three windows of the EFTOS Monitor. In window 1 ("Global view"), the `visual` column contains graphical hyperlinks pointing to second-level information about the corresponding processing node at the same row. `Configuration` is the DIR net-role. `Status` may be one of the following values: `OK`, `Faulty`, `Isolated`, `Recovering`, and `Killed`. Some minor information is also displayed at the bottom of the page. The right-hand hypertext (window "Node-specific information") is the result of "clicking" on the top circular icon and enumerates the actions that have just taken place on node 0, fresher-to-older. The elapsed time (in seconds) corresponding to each event is displayed. Underlined sentences may be further expanded by clicking on them e.g., the bottom-left image reports about action number 115 of the hypertext.

Fig. 1, a control path may be drawn starting from the user at his/her browser, then crossing a CGI script, and eventually reaching the user application. It is therefore fairly possible to add a layer to the hierarchy of HTML pages dynamically created by the intermediate modules so that the user may freely choose among a certain set of malicious actions to bring against an EFTOS application, including for instance:

- an integer division-by-zero,
- a segmentation violation,
- a link failure,
- rebooting a processing node,
- killing a thread.

These requests would then reach a CGI script, be translated in appropriate system- or application-level

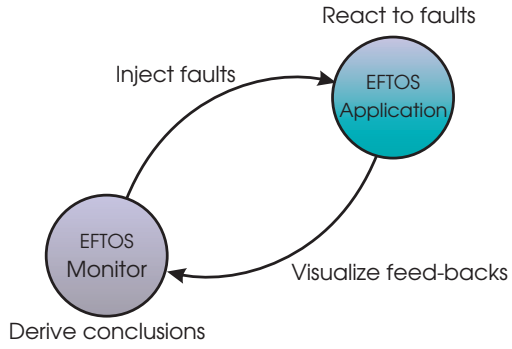

Figure 3: The recursive loop of fault injection and monitoring.

actions, which would then be executed or turned into fault-injection requests to be fulfilled by the DIR Manager. As an example of system-level action, the CGI script may directly execute a system command to reboot one node in the CC system. As of application-level actions, the Manager may for instance ask the trap handler tool to trigger a specific signal like `SIGSEGV` (segmentation violation) on a certain thread; or it may request a watchdog timer tool on a particular node to behave like if it had detected a time-out.

As a direct consequence of the injection of these faults, a number of detection, isolation, and recovery actions will take place on the system according to the EFTOS-based fault-tolerance strategies adopted by the designer in his/her application. These actions will then be reported in the snapshot files and displayed by the Monitor. This process, summarized in Fig. 3, may be modeled after a recursive loop like follows:

```
do {
    Inject fault;
    Observe feed-back;
    Derive conclusions;
    Correct the fault tolerance model;
} while (model is unsatisfying).
```

In our opinion this procedure should result in an extremely useful tool for rapidly assessing a design, trying alternative fault-tolerance strategies, and overloading the system with malicious attacks aiming at verifying its resilience, with a quick and meaningful feed-back from the system.

6 Conclusions

We presented the current state of development of a distributed application for monitoring the fault tolerance aspects of an embedded parallel application and for interactively injecting faults into it. The overall system makes up an integrated environment in which they cyclically evolve: the application, a sophisticated graphical rendering of the results, and real-time interactions such that the researcher is made able to verify the hypothesis he/she is formulating about the system.

The design choice to adopt low-cost, off-the-shelf components for hypermedia rendering revealed to be cost-effective, to speed up the development process, to match the design requirements, and to point at more ambitious capabilities and features. In particular, the use of a World-Wide Web browser as hypermedia renderer paves the way for future client-based extensions based on JavaScript or Java [7], and lets our application inherit the benefits of the volcanic evolutions of the HTML languages, the HTTP protocol, multimedia capabilities of the browsers, and so on.

The high degree of openness proven by this heterogeneous application basing itself on uniform communication mechanisms and standardized access interfaces guarantees portability and makes it also a good starting point towards the development of similarly structured applications ranging from remote equipment control to hypermedia multi-user environments.

We are currently using our Monitor during the development of the new versions of the EFTOS fault tolerance framework. The deeper insight that we have gained from it on the run-time aspects of our applications has turned into an invaluable tool to speed up our development phases.

Acknowledgements

This project is partly sponsored by an FWO Krediet aan Navorsers, by the Esprit-IV Project 21012 EFTOS, and by COF/96/11. Vincenzo De Florio is on leave from Tecnopolis CSATA Novus Ortus. Geert Deconinck is a Postdoctoral Fellow of the Fund for Scientific Research - Flanders (Belgium) (F.W.O.). Rudy Lauwereins is a Senior Research Associate of F.W.O.

References

[1] Anonymous, *Parsytec CC Series—Cognitive Computing,* Parsytec GmbH, Aachen, 1996.

[2] Anonymous, "Embedded Parix Programmer's Guide," *Parsytec CC Series Hardware Documentation,* Parsytec GmbH, Aachen, 1996.

[3] T.J. Berners-Lee, L. Masinter, and M. McCahill, "Uniform Resource Locators (URL)," Request for Comments Vol. 1738, Network Working Group, Dec. 1994.

[4] T.J. Berners-Lee, R. Cailliau, J.-F. Groff, and B. Pollermann, "World-Wide Web: the Information Universe," *Electronic Networking: Research, Applications and Policy*, Vol. 2, No.1, pp.52–58, Meckler, Westport, 1992.

[5] T.J. Berners-Lee and D. Connolly, "Hypertext Markup Language — 2.0," Request for Comments Vol. 1866, Network Working Group, Nov. 1995.

[6] T.J. Berners-Lee, R. Fielding, and H. Frystyk, "Hypertext Transfer Protocol — HTTP/1.0," Request for Comments Vol. 1945, Network Working Group, May 1996.

[7] M. Campione and K. Walrath, *The Java Tutorial — Object-Oriented Programming for the Internet,* Addison-Wesley, New York, 1996.

[8] N. Carriero and D. Gelernter, "How to write parallel programs: a guide to the perplexed," *ACM Comp. Surv.* Vol. 21, pp. 323–357, 1989.

[9] D.E. Comer and D.L. Stevens, *Internetworking with TCP/IP, Vol. 3: Client-Server Programming and Applications,* Prentice-Hall, Englewood-Cliffs, 1993.

[10] G. Deconinck, V. De Florio, R. Lauwereins, and T. Varvarigou, "EFTOS: A Software Framework for More Dependable Embedded HPC Applications," Proc. of the Third Int. Euro-Par Conference, *Lecture Notes in Computer Science,* Vol. 1330, pp.1363–1368, Springer, Berlin, 1997.

[11] V. De Florio, "L'Azienda Virtuale Mudhoney," *Internet News* no.7, Tecniche Nuove, Milano, 1995.

[12] V. De Florio, "Oltre la CGI: lo Standard Common Client Interface," *DEV.* no.29, Infomedia, Pescara, 1996.

[13] R. Fielding, R., U.C. Irvine, J. Gettys, J. Mogul, H. Frystyk, and T.J. Berners-Lee, "Hypertext Transfer Protocol — HTTP/1.1," Request for Comments Vol. 2068, Network Working Group, Jan. 1997.

[14] B.W. Johnson, *Design and Analysis of Fault-Tolerant Digital Systems,* Addison-Wesley, New York, 1989.

[15] E.E. Kim, *CGI Developer's Guide,* SAMS.NET, 1996.

[16] J. Zawinski, "Remote Control of UNIX Netscape," URL: `http://home.netscape.com/newsref/std/xremote.html`, Netscape Communications Corp., 1994.